



---

**BRI-FY13-Ultra-Scalable Algorithms for Large-scale Uncertainty Quantification**

**Omar Ghattas**  
**UNIVERSITY OF TEXAS AT AUSTIN**

---

**03/04/2016**  
**Final Report**

DISTRIBUTION A: Distribution approved for public release.

Air Force Research Laboratory  
AF Office Of Scientific Research (AFOSR)/ RTA2  
Arlington, Virginia 22203  
Air Force Materiel Command

# REPORT DOCUMENTATION PAGE

Form Approved  
OMB No. 0704-0188

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. **PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.**

<b>1. REPORT DATE</b> (DD-MM-YYYY) 29-2-2016		<b>2. REPORT TYPE</b> Final report		<b>3. DATES COVERED</b> (From — To) 30 Sept 2012 – 30 Nov 2015	
<b>4. TITLE AND SUBTITLE</b>  Ultra-Scalable Algorithms for Large-Scale Uncertainty Quantification in Inverse Wave Propagation				<b>5a. CONTRACT NUMBER</b>	
				<b>5b. GRANT NUMBER</b> FA9550-12-1-0484	
				<b>5c. PROGRAM ELEMENT NUMBER</b>	
<b>6. AUTHOR(S)</b>  George Biros, Leszek Demkowicz, Omar Ghattas, Jay Gopalakrishnan				<b>5d. PROJECT NUMBER</b>	
				<b>5e. TASK NUMBER</b>	
				<b>5f. WORK UNIT NUMBER</b>	
<b>7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)</b> Institute for Computational Engineering & Sciences The University of Texas at Austin Stop C0200 Austin, TX 78712-1229				<b>8. PERFORMING ORGANIZATION REPORT NUMBER</b>	
<b>9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)</b>  Air Force Office of Scientific Research Computational Mathematics Program 875 North Randolph Street, Suite 325 Arlington, VA 22203				<b>10. SPONSOR/MONITOR'S ACRONYM(S)</b>  AFOSR	
				<b>11. SPONSOR/MONITOR'S REPORT NUMBER(S)</b>	
<b>12. DISTRIBUTION / AVAILABILITY STATEMENT</b>  Approved for public release. DISTRIBUTION A					
<b>13. SUPPLEMENTARY NOTES</b>					
<b>14. ABSTRACT</b> The overall aim of this project was to develop scalable algorithms for the inverse problem of inferring, with associated uncertainty, the heterogeneity of a medium or shape of a scatterer from reflected/transmitted waves (acoustic, elastic, electromagnetic) at very large scale. The resulting Bayesian wave inverse propagation problem has been intractable using contemporary algorithms. Research was conducted under three complementary subprojects. The first subproject (led by O. Ghattas) focused on scalable algorithms for large-scale Bayesian inverse problems governed by time domain wave propagation. The second subproject (led by G. Biros) focused on fast algorithms for inverse scattering and uncertainty quantification based on volume integral equation formulations for the inverse medium problem. The third subproject (led by L. Demkowicz and J. Gopalakrishnan) focused on new, highly efficient discretizations for wave propagation in the form of the discontinuous Petrov Galerkin (DPG) method and associated solvers. Results and conclusions in each sub-project area are discussed in separate sections of the report.					
<b>15. SUBJECT TERMS</b> Acoustic, elastic, and electromagnetic wave propagation; discontinuous Petrov Galerkin method; volume integral equations; fast multipole method; FFT; inverse medium and inverse shape scattering; Bayesian inverse problems; uncertainty quantification; parallel algorithms; scalable algorithms; high performance computing					
<b>16. SECURITY CLASSIFICATION OF:</b>			<b>17. LIMITATION OF ABSTRACT</b>	<b>18. NUMBER OF PAGES</b>  13	<b>19a. NAME OF RESPONSIBLE PERSON</b> Omar Ghattas
<b>a. REPORT</b>  UU	<b>b. ABSTRACT</b>  UU	<b>c. THIS PAGE</b>  UU			<b>19b. TELEPHONE NUMBER</b> (include area code) 512-232-4304

## Final Report for AFOSR grant FA9550-12-1-0484

### Ultra-Scalable Algorithms for Large-Scale Uncertainty Quantification in Inverse Wave Propagation

PIs: G. Biros, L. Demkowicz, O. Ghattas (Lead PI), J. Gopalakrishnan

This overall aim of this project was to address the inverse problem of inferring, with associated uncertainty, the heterogeneity of a medium or shape of a scatterer from reflected/transmitted waves (acoustic, elastic, electromagnetic) at very large scale. The resulting Bayesian wave inverse propagation problem has been intractable using contemporary algorithms. Research was conducted along three complementary subprojects. The first subproject (led by O. Ghattas) focused on scalable algorithms for large-scale Bayesian inverse problems governed by time domain wave propagation. The second subproject (led by G. Biros) focused on fast algorithms for inverse scattering and uncertainty quantification based on volume integral equation formulations for the inverse medium problem. The third subproject (led by L. Demkowicz and J. Gopalakrishnan) focused on new, highly efficient discretizations for wave propagation in the form of the discontinuous Petrov Galerkin (DPG) method and associated solvers. Each subproject is described below.

#### 1. Scalable algorithms for large-scale Bayesian inverse medium and shape problems governed by time domain wave propagation

This component of the project addressed the problem of Bayesian inverse problems governed by time-domain wave propagation (acoustic, elastic, and electromagnetic). Results were obtained along the following lines:

- **Extreme-scale UQ for Bayesian inverse wave propagation.** We developed parallel algorithms and implementations for extreme scale inverse problems governed by the acoustic/elastic wave equation in the Bayesian inference framework: given data and model uncertainties, find the pdf describing parameter uncertainties. To overcome the curse of dimensionality of conventional methods, we exploit the fact that the data are typically informative about low-dimensional manifolds of parameter space. This leads to a low rank approximation of the prior-preconditioned Hessian of the negative log likelihood, evaluated at the maximum a posteriori (MAP) point and effected via matrix-free randomized SVD, in conjunction with a Sherman Morrison Woodbury inverse, to arrive at a Gaussianized approximation of the posterior covariance [17]. We obtain a method that scales independent of the forward problem dimension, the uncertain parameter dimension, the data dimension, and the number of cores. This approximation is exact for a linear parameter-to-observable map (modulo controllable error in the randomized SVD), and forms the basis for a locally-adaptive Gaussian proposal density in a Metropolis Hastings MCMC method [53]. The largest problem solved had a million uncertain parameters with 630 million DOF, on up to 262K cores, for which a factor of 2000 reduction in parameter dimension was achieved [10]. This remains the largest Bayesian inverse wave propagation ever solved.
- **Fast solvers for Bayesian priors for inverse wave propagation in layered media.** One popular choice of prior covariance operator for Bayesian inverse problems is the inverse of power of a Laplacian-like elliptic differential operator. The motivation for this choice is that the operator allows heterogeneous and anisotropic control over correlation lengths and variance, and is of trace class, leading to a well-posed Bayesian inverse problem. However, this means that whenever the prior is manipulated (in practice, thousands of times at a minimum), an elliptic PDE must be

solved. Thus in order to achieve scalability in Bayesian inversion, we need a multigrid solver that can scale to extreme core counts. We have designed such a multigrid solver using a combination of algebraic and geometric ideas [59], extending to high-order discretization on complex geometries [61], and scaling in the largest cases up to 1.5 million cores with 602 billion unknowns [57].

- **A mathematical and computational framework for infinite-dimensional Bayesian inverse wave propagation.** The mathematical and computational basis for the extreme scale algorithms described above has been presented in a series of papers addressing infinite-dimensional Bayesian inverse problems. We began with the ideas proposed by Stuart (*Acta Numerica*, 2010), and incorporated a number of components aimed at ensuring a convergent discretization of the underlying infinite-dimensional inverse problem. The framework additionally incorporated algorithms for manipulating the prior, constructing a low rank approximation of the data-informed component of the posterior covariance operator, and exploring the posterior that together ensure scalability of the entire framework to very high parameter dimensions. The framework was established first for a linearized parameter-to-observable (p2o) map [17], and then extended to a fully nonlinear p2o map by using the resulting locally-adapted Gaussianized posterior approximation as a proposal for a Metropolis-Hastings Markov Chain Monte Carlo method [53]. Finally, we analyzed a model inverse scattering problem (specifically, inverse shape acoustic scattering) and showed well-posedness of the Bayesian formulation in infinite dimensions, as well as convergence of the finite-dimensional approximation (of both the shape and the state) to the infinite-dimensional posterior measure, in which convergence rates of the finite-dimensional inverse problem are inherited from those of both the prior (on the shape) and the forward wave propagation problem [15]. We have employed this framework in inverse wave propagation problems involving subsurface mapping of realistic media [66].
- **Analysis of the Hessian for 3D inverse scattering.** The scalability of the Bayesian inversion methodology described above is intimately tied to the low rank approximation of the (prior-preconditioned) Hessian of the negative log of the likelihood (i.e., the Hessian of the weighted data misfit functional). This low rank approximation is motivated by the fact that for infinite-dimensional inverse problems, the data typically inform a low-dimensional manifold of parameter space (hence the ill-posedness of the unregularized inverse problem), leading to a compact data misfit Hessian operator. For many problems we do numerically observe a rapid decrease of eigenvalues of this operator, permitting up to three orders of magnitude (implicit) dimension reduction. Continuing work conducted under our previous AFOSR grant, which theoretically verified the compactness of the data misfit Hessian for 2D inverse shape [12] and inverse medium [13] acoustic scattering, we extended the theoretical analysis to the more difficult case of 3D scattering with electromagnetic waves and showed that here the Hessian is also a compact operator [14].
- **Discretely exact derivatives for hyperbolic PDE-constrained optimization problems discretized by the discontinuous Galerkin method.** The adjoint wave equation is a critical component for efficiently computing the gradient of log posterior (to define the MAP point) and its Hessian (to construct a posterior covariance approximation). However a technical point arises when the wave equation is discretized by the discontinuous Galerkin method: should the continuous adjoint wave equation be derived and then discretized by DG (in which case the resulting discretized gradient may not be consistent with the log posterior functional)? Or should we first discretize the forward wave equation by DG and then derive an adjoint (in which case the adjoint wave equation may not correspond to a DG discretization of the forward equation)? We have analyzed these two alternatives [64] and have shown that the gradient in the former approach is inconsistent with the discretized gradient, leading to a possible lack of convergence of gradient-

based optimization methods. Moreover, we have shown that the discrete adjoint equation inherits a natural DG discretization from the discretization of the forward wave equation, and the resulting gradient expressions have to take into account additional contributions from element faces in order to be discretely exact and thus lead to the correct gradient for numerical optimization purposes. Our Bayesian inversion framework employs the results of this paper for DG discretization of the adjoint wave equation and the resulting gradient.

- **Besov priors for preserving medium discontinuities in Bayesian inverse problems.** A critical issue when solving Bayesian inverse wave propagation problems involving inhomogeneous media with material property jumps (such as layered media) is the specification of a prior covariance operator that preserves jumps. For deterministic inverse wave propagation problems, this can be achieved by total variation (TV) regularization [1, 2, 31]. Unfortunately, TV does not converge in the limit of finer discretization. Recently, Besov space priors (which involve  $\ell_1$ -regularized wavelet coefficients of the medium property field) were proposed as a means of retaining discretization invariance as well as edge preservation. In our work we have introduced a fast solver for such priors that is substantially faster than existing approaches (split Bregman and interior path following primal-dual methods) [16].
- **Fast optimization-based MCMC sampling methods for posteriors for Bayesian inverse wave propagation problems.** We developed a so-called randomized maximum a posteriori (rMAP) method for generating approximate samples of posteriors in high dimensional Bayesian inverse problems governed by large-scale forward problems, with particular application to wave propagation [63]. The rMAP approach is derived based on: 1) casting the problem of computing the MAP point as a stochastic optimization problem; 2) interchanging optimization and expectation; and 3) approximating the expectation with a Monte Carlo method. For a specific randomized data and prior mean, rMAP reduces to the maximum likelihood approach (RML). It can also be viewed as an iterative stochastic Newton method. An analysis of the convergence of the rMAP samples was carried out for both linear and nonlinear inverse problems. Each rMAP sample requires solution of a PDE-constrained optimization problem; to solve these problems, we employed a state-of-the-art trust region inexact Newton conjugate gradient method with sensitivity-based warm starts. An approximate Metropolization approach is presented to reduce the bias in rMAP samples. This method can be thought of as an extension of our previously-developed stochastic Newton method [51] (which employs a Gaussian proposal based on a covariance operator taken to be the inverse of the local Hessian) to a non-Gaussian proposal based on a nonlinear trajectory in parameter space; indeed the two are equivalent for a linear parameter-to-observable map. Numerical results indicated the potential of the rMAP approach in posterior sampling of nonlinear Bayesian inverse wave propagation problems in high dimensions.
- **Optimal source compression for Bayesian inverse wave propagation problems based on optimal experimental design.** A major challenge for inverse wave propagation problems is in the common situation when there are multiple sources to interrogate the medium, rather than one. In this case, the forward (and adjoint) wave equation must be solved multiple times per gradient or Hessian-vector evaluation, once for each source. This means that for many industrial settings, hundreds or more wave equations will need to be solved at each inversion iteration (or MCMC sample point). Clearly, these sources do not all yield independent information on the medium. Can they be collapsed into a handful of “meta-sources”? Inspired by our work on optimal experimental design for Bayesian inverse problems [3, 4], in the present project we formulated an approach to this problem based on optimal experimental design [22]. That is, find a small number of optimal linear combinations of the sources such that the medium is recovered with the least

uncertainty. For an uncertainty measure, we adopted the A-optimal design criterion, i.e., the trace of the inverse of the Hessian (of the log posterior) evaluated at the MAP point. This results in a PDE-constrained optimization problem that is constrained by the vanishing of the gradient (along with forward/adjoint wave equations to define the gradient), as well as linear systems with Hessian operators that arise in the trace estimation (along with incremental forward/adjoint wave equations to determine the Hessian action).

The work described above was carried out by research associates/scientists who have now moved on to faculty positions: Alen Alexanderian (NC State), Tan Bui-Thanh (UT-Austin), Carsten Burstedde (University of Bonn), Noemi Petra (UC Merced), Georg Stalder (NYU), Hari Sundar (University of Utah), and Lucas Wilcox (Naval Postgraduate School).

## 2. Fast algorithms for inverse scattering and uncertainty quantification based on volume integral equation formulations for inverse medium problems

We worked on fast algorithms for inverse scattering and uncertainty quantification based on volume integral equation formulations for the inverse medium problem. The fast solvers include forward and adjoint scattering solvers, parallel algorithms for Hessian approximations, and fundamental algorithms for kernel methods.

Our goal is to develop algorithms that allow the efficient solution of forward, inverse, and UQ problems for the frequency-domain formulation of scalar and vector scattering problems in inhomogeneous media (that is, media for which the refractive index varies in space). Funding from this award led to 14 publications in premier peer-reviewed journals and conferences and three software libraries (PVFMM, AccFFT, LIBASKIT). It has partially supported five PhD students (Chenhan Yu, Dhairya Malhotra, Amir Gholami, Bo Xiao, and Keith Kelly) and three postdoctoral scientists (Hari Sundar, Brian Quaife, and Bill March), the first two of whom have assumed assistant professor positions (at the University of Utah and Florida State University). At the ACM/IEEE SC'15 conference, Dhairya Malhotra received the ACM George Michael HPC Fellowship Award for his work on the fast multipole method and Amir Gholami received the Gold Prize in the ACM Student Research Competition for his work on the fast Fourier transform.

**Fast solvers for forward, adjoint, and Hessian systems.** Volume integral equations enable the solution of scattering problems with high-order accuracy, have negligible dispersion errors, capture radiation conditions exactly, and offer unprecedented algorithmic and parallel scalability. We have developed two solvers for volume integral equations, one accelerated by the *fast multipole method* (FMM) and one accelerated by the *fast Fourier transform* (FFT). The FMM-based solver allows highly adaptive discretizations. However, it is limited to low and medium frequency problems (roughly speaking, up to 100 wavelengths). The FFT-based one enables the solution of scattering problems with arbitrarily high frequencies but it uses regular grids. Both methods scale to at least  $\mathcal{O}(10^5)$  cores. For both approaches, we have developed open source software and have made it freely available. We give more details on our technical contributions below.

- **Highlights for FMM.** Our PVFMM (parallel kernel independent fast multipole method for volume potentials) can be used to construct spatially-adaptive solvers for the Poisson, Stokes, and low-frequency Helmholtz problems. Conventional N-body methods apply to discrete particle interactions. With volume potentials, we replace the sums with volume integrals. We use high-order piecewise Chebyshev polynomials and an octree data structure to represent the input and output fields, enable spectrally accurate approximation of the near-field, and the kernel independent FMM (KIFMM) for the far-field approximation. For distributed-memory parallelism, we use space filling curves, locally essential trees, and a hypercube-like collective communication scheme. Our

PVFMM can achieve about 600 GF/s of double-precision performance on a single node. Our largest run on the *Stampede* system at the Texas Advanced Computing Center took 3.5s on 16K cores for a problem with  $18\text{E}+9$  unknowns for a highly nonuniform scattering field (corresponding to an effective resolution exceeding  $3\text{E}+23$  unknowns since we used 23 levels in our octree). The code is publicly available at <http://www.pvfmm.org>. Related publications include [6, 43, 44, 55, 62].

- **Highlights for FFT.** Despite the large amount of work on FFTs, we have shown that significant speedups can be achieved for distributed transforms. AccFFT extends existing FFT libraries for x86 architectures (CPUs) and CUDA-enabled Graphics Processing Units (GPUs) to distributed memory clusters using the Message Passing Interface (MPI). Our library uses specifically optimized all-to-all communication algorithms to efficiently perform the communication phase of the distributed FFT algorithm. We tested our library on the *Maverick* and *Stampede* systems at TACC and on the *Titan* system at Oak Ridge National Laboratory. The library was tested on up to 131K cores and 4,096 GPUs of *Titan*, and on up to 16K cores of *Stampede*. The library is available at <http://www.accfft.org>. The main publication for the FFT is [33]; we have also compared FFT, FMM, and multigrid for solving elliptic PDEs at [34].

As mentioned above, both FMM-based and FFT-based solvers have their limitations. Our immediate goal is to combine the two techniques in a single solver, combining the benefits of both without the limitations. The new solver should be efficient for very high-frequency problems while supporting spatially non-uniform discretizations.

These solvers have been incorporated into fast algorithms for Hessian approximations. The adjoint for the FFT-accelerated scheme is simple, but the adjoint for the FMM is more involved. Also, we have implemented our FaIMS method [20] (developed with previous funding from AFOSR) in parallel, integrated it with volume integrals (not just charges), and combined it with the Elemental library [54] to enable fast and scalable randomized linear algebra. We have also been working on domain decomposition preconditioners.

**Algorithms for uncertainty quantification.** We have been exploring fundamental algorithms that enable several new approaches in uncertainty quantification. The main motivation is to find an algorithmic way to encapsulate priors given as kernel densities and are problem specific—instead of exclusively using smoothness priors. However, using kernel densities as priors poses significant computational challenges. In a series of papers, we developed *hierarchical matrix* technology that is applicable to kernel density estimation but also to fast approximations for the Hessian operators in inverse scattering problems. Our goal is a method that exhibits algorithmic and parallel scalability for inverting kernel and Hessian matrices. We require only the ability to compute an entry of the target matrix in  $\mathcal{O}(\log N)$  time, where  $N$  is the number of unknowns. We developed a new Approximate Skeletonization Kernel Independent Treecode (ASKIT) that builds technology for a special class of hierarchical matrices. The ASKIT library is available at <http://padas.ices.utexas.edu/libaskit>. In a series of papers we introduced the new technology [45–50] enabling fast matrix-vector multiplies in  $\mathcal{O}(N)$  time. Furthermore in our most recent work [65], we developed a *fast direct solver* for hierarchical matrices that can be used to precondition the prior in the Hessian. The factorization requires  $\mathcal{O}(N \log^2 N)$  work. To the best of our knowledge, this direct solver represents the state-of-the-art. It can be used with high-dimensional data for a large variety of kernels and is the only algorithm that supports shared and distributed memory parallelism. Our scheme does not assume symmetry of the kernel, global low-rank structure, sparsity, or any other property other than, up to a sparse matrix correction, the off-diagonal blocks admit a low-rank approximation.



### 3. New discontinuous Petrov Galerkin methods for wave propagation problems

This project focused on the development of efficient high quality wave simulators using discontinuous Petrov Galerkin methods (DPG), invented previously by L. Demkowicz and J. Gopalakrishnan. Publications directly or indirectly influenced by this project are [5, 7, 9, 11, 18, 19, 21, 24–30, 32, 35–40, 42, 52, 58]. Highlights are listed below.

- **Foundational work on mathematical error analysis** of DPG methods is now complete. Conditions for a priori and a posteriori error analyses of the abstract DPG method were found [18, 19, 39]. The results generalize essentially our earlier work on the subject and are as follows. First, any linear boundary-value or initial-boundary-value problem, wave propagation problems included, admits many variational formulations employing different energy settings and implying ultimate convergence in different norms [19, 24, 42]. These problems are *simultaneously* well- or ill-posed. Second, in each of these formulations, the standard test spaces can be replaced with *broken test spaces* resulting in additional unknowns defined on mesh skeletons (*traces*) but enabling the application of the DPG technology. This does not mean that all formulations are “equal.” Being a minimum-residual method, DPG method always delivers a positive definite Hermitian stiffness matrix but its spectral properties vary dramatically between the formulations. In this context, the so-called *ultra-weak* formulation stands out. It is *robust* (uniformly stable with respect the frequency), and it delivers the best conditioned stiffness matrix among the different DPG formulations.

We distinguish between the *ideal* and *practical* DPG methods. In the analysis of the ideal DPG method, we assume that the optimal test functions are computed exactly. In the analysis of the practical DPG method, we account for the approximation of optimal test functions by constructing appropriate Fortin operators [19, 39, 58]. The Fortin operators are also crucial in the analysis of a-posteriori error estimation [18].

- When using the DPG method, **adaptivity was found to be robust without any preasymptotic instabilities** [21, 29, 30]. This property should be strongly contrasted with the standard Galerkin method which is only *asymptotically stable*. In practice this means that, with standard Galerkin, we have to start with a mesh that not only satisfies the Nyquist criterion (resolves the wave number) but also controls the phase error (the so-called *pollution error*; high order elements are probably the best methodology here). Outside of the asymptotic regime, standard Galerkin is completely unreliable, in particular, a-posteriori error estimates do not work, and adaptivity is disabled. Contrary to the standard Galerkin, minimum residual methods, DPG methodology included, come with an a-posteriori error estimate (the residual) built in, and enable adaptivity from day one, starting with very coarse meshes. The adaptive DPG technology is attractive for very high frequency problems ( $> 300$  wavelengths in 2D) with “localized solutions” (beams) where it stands a chance to beat standard Galerkin methods. We are in the process of developing a special solver that integrates adaptivity with domain decomposition methods.
- **Phase errors reduce** with the right choice of test norm in DPG methods. Reduction of dissipation was proved using a numerical dispersion analysis in [35, 38]. The dispersion analysis confirms the difference between different variational formulations. The DPG method based on the strong formulation reduces to the classical First Order System Least Squares (FOSLS) method and displays the strongest dissipation while the ultra-weak DPG method has the least dissipation properties. The dissipation can be further reduced by scaling the  $L^2$ -terms in the adjoint graph norm [38]. Such a scaling changes effectively the norm in which traces are measured, see also [19].



- **DPG method works when standard Galerkin method might not work**, as shown by studies on metamaterials. The DPG method was applied to simulate cloaking and was found to be very efficient with the right choice of the test norm [28]. If the original problem is well-posed, i.e. the corresponding sesquilinear form satisfies the inf-sup condition, the DPG method guarantees reproducing the continuous stability at the discrete level. This is critical for wave propagation problems in metamaterials which are well-posed (satisfy the inf-sup condition) but do not satisfy the criteria for Mikhlin's theory of asymptotic stability (applicable to standard Galerkin and standard wave propagation problems). For such classes of problems the standard Galerkin may fail to converge.
- A **frequency-independent Schwarz preconditioner** for DPG methods was developed [40]. We found that a careful choice of overlapping blocks within a multiplicative Schwarz algorithm, applied to the DPG system, with no coarse solve, provides a preconditioner for solving the DPG system for waves. Its performance is, as expected from similar known results, independent of the polynomial degree  $p$ . What was unexpected was a pleasant observation that the condition number of the preconditioned system is independent of the wavenumber (in addition to  $p$ ). Together with the solvability of DPG methods on any mesh, no matter how coarse in relation to the wavenumber, this allows us to design a robust preconditioned solution strategy for wave simulations.
- DPG methods have been successfully **applied and work for many formulations of Maxwell equations** [19]. We showed that the same problem admits different variational formulations leading to different FE discretizations and different types of convergence [24]. We proved that **the stability of one implies the stability of the others**. Maxwell equations constitute an important particular example of the general theory discussed above. In particular, one of the main challenges here was the construction of an appropriate Fortin operator.
- Finally, the DPG method has been successfully applied to **inverse seismic tomography problems** [8] proving to be an attractive alternative to standard high order Galerkin method in context of multifrequency inversion schemes.
- **Software:** The grant supported the development of a Fortran90 codebase implementing DPG methods within the existing framework of the *hp* FEM software [23], a C++ codebase available publicly as a GIT repository [40, 41] (designed as a shared library add-on to a popular open source package NGSolve), and the initial development of a stand alone C++ codebase solely for DPG methods called CAMELIA [56]. In particular, in the course of this project, we developed a special library for computing *orientation embedded* hierarchical shape functions of arbitrary order, elements of all shapes (tetrahedra, hexahedra, prism *and pyramids* in 3D) and the spaces forming the first Nedélec exact sequence [32]. The library (9k of Fortran 90 code) is in the public domain and can be used in any higher order FE code.
- The grant supported these completed or upcoming **Ph.D. dissertations/students**:  
 Austin: J. Bramwell (2013), S. Nagaraj (2017), S. Petrides (2017)  
 Portland: N. Olivares (2016), A. Harb (2016), P. Sepulveda (2017).

## References

- [1] V. AKÇELİK, J. BIELAK, G. BIROS, I. EPANOMERITAKIS, A. FERNANDEZ, O. GHATTAS, E. J. KIM, J. LOPEZ, D. R. O'HALLARON, T. TU, AND J. URBANIC, *High resolution forward and inverse earthquake modeling on terascale computers*, in SC03: Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis, ACM/IEEE, 2003. Gordon Bell Prize for Special Achievement.
- [2] V. AKÇELİK, G. BIROS, AND O. GHATTAS, *Parallel multiscale Gauss-Newton-Krylov methods for inverse wave propagation*, in Proceedings of IEEE/ACM SC2002 Conference, Baltimore, MD, Nov. 2002. SC2002 Best Technical Paper Award.
- [3] A. ALEXANDERIAN, N. PETRA, G. STADLER, AND O. GHATTAS, *A-optimal design of experiments for infinite-dimensional Bayesian linear inverse problems with regularized  $\ell_0$ -sparsification*, SIAM Journal on Scientific Computing, 36 (2014), pp. A2122–A2148.
- [4] A. ALEXANDERIAN, N. PETRA, G. STADLER, AND O. GHATTAS, *A fast and scalable method for A-optimal design of experiments for infinite-dimensional Bayesian nonlinear inverse problems*, SIAM Journal on Scientific Computing, 38 (2016), pp. A243–A272.
- [5] D. M. AMBROSE, J. GOPALAKRISHNAN, S. MOSKOW, AND S. ROME, *Scattering of electromagnetic waves by thin high contrast dielectrics II: Asymptotics of the electric field and a method for inversion*, Preprint, (2016).
- [6] G. BIROS AND D. MALHOTRA, *PVFMM: A parallel kernel independent FMM for particle and volume potentials*, Communications in Computational Physics, 18 (2015), pp. 808–830.
- [7] T. BOUMA, J. GOPALAKRISHNAN, AND A. HARB, *Convergence rates of the DPG method with reduced test space degree*, Computers and Mathematics with Applications, 68 (2014), pp. 1550–1561.
- [8] J. BRAMWELL, *A Discontinuous Petrov-Galerkin Method for Seismic Tomography Problems*, PhD thesis, University of Texas at Austin, 2013. supervisors: L. Demkowicz and O. Ghattas.
- [9] J. BRAMWELL, L. DEMKOWICZ, J. GOPALAKRISHNAN, AND W. QIU, *A locking-free hp DPG method for linear elasticity with symmetric stresses*, Numer. Math., 122 (2012), pp. 671–707.
- [10] T. BUI-THANH, C. BURSTEDDE, O. GHATTAS, J. MARTIN, G. STADLER, AND L. C. WILCOX, *Extreme-scale UQ for Bayesian inverse problems governed by PDEs*, in SC12: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, 2012. Gordon Bell Prize finalist.
- [11] T. BUI-THANH, L. DEMKOWICZ, AND O. GHATTAS, *A unified discontinuous Petrov-Galerkin method and its analysis for Friedrichs' systems*, SIAM Journal on Numerical Analysis, 51 (2013), pp. 1933–1958.
- [12] T. BUI-THANH AND O. GHATTAS, *Analysis of the Hessian for inverse scattering problems. Part I: Inverse shape scattering of acoustic waves*, Inverse Problems, 28 (2012), p. 055001.
- [13] ———, *Analysis of the Hessian for inverse scattering problems. Part II: Inverse medium scattering of acoustic waves*, Inverse Problems, 28 (2012), p. 055002.

- [14] ———, *Analysis of the Hessian for inverse scattering problems. Part III: Inverse medium scattering of electromagnetic waves*, Inverse Problems and Imaging, 7 (2013), pp. 1139–1155.
- [15] ———, *An analysis of infinite dimensional Bayesian inverse shape acoustic scattering and its numerical approximation*, SIAM Journal of Uncertainty Quantification, 2 (2014), pp. 203–222.
- [16] ———, *A scalable MAP solver for Bayesian inverse problems with Besov priors*, Inverse Problems and Imaging, 9 (2015), pp. 27–54.
- [17] T. BUI-THANH, O. GHATTAS, J. MARTIN, AND G. STADLER, *A computational framework for infinite-dimensional Bayesian inverse problems Part I: The linearized case, with application to global seismic inversion*, SIAM Journal on Scientific Computing, 35 (2013), pp. A2494–A2523.
- [18] C. CARSTENSEN, L. DEMKOWICZ, AND J. GOPALAKRISHNAN, *A posteriori error control for DPG methods*, SIAM J Numer. Anal., 52 (2014), pp. 1335–1353.
- [19] C. CARSTENSEN, L. DEMKOWICZ, AND J. GOPALAKRISHNAN, *Breaking spaces and forms for the dpg method and applications including maxwell equations*, Preprint arXiv:1507.05428, (2015).
- [20] S. CHAILLAT AND G. BIROS, *FaIMS: A fast algorithm for the inverse medium problem with multiple frequencies and multiple sources for the scalar helmholtz equation*, Journal of Computational Physics, 231 (2012), pp. 4403–4421.
- [21] J. CHAN, N. HEUER, T. BUI-THANH, AND L. DEMKOWICZ, *A robust DPG method for convection-dominated diffusion problems II: adjoint boundary conditions and mesh-dependent test norms*, Comput. Math. Appl., 67 (2014), pp. 771–795.
- [22] B. CRESTEL, A. ALEXANDERIAN, G. STADLER, AND O. GHATTAS, *A-optimal encoding weights for nonlinear inverse problems, with application to the Helmholtz inverse problem*, Submitted, (2016).
- [23] L. DEMKOWICZ, *Computing with hp-adaptive finite elements. Vol. 1*, Chapman & Hall/CRC Applied Mathematics and Nonlinear Science Series, Chapman & Hall/CRC, Boca Raton, FL, 2007.
- [24] L. DEMKOWICZ, *Various variational formulations and closed range theorem*, ICES Report, (2015).
- [25] L. DEMKOWICZ AND J. GOPALAKRISHNAN, *An overview of the discontinuous Petrov Galerkin method*, in Recent Developments in Discontinuous Galerkin Finite Element Methods for Partial Differential Equations: 2012 John H Barret Memorial Lectures, X. Feng, O. Karakashian, and Y. Xing, eds., vol. 157 of The IMA Volumes in Mathematics and its Applications, Institute for Mathematics and its Applications, Minneapolis, Springer, 2013, pp. 149–180.
- [26] L. DEMKOWICZ AND J. GOPALAKRISHNAN, *A primal DPG method without a first-order reformulation*, Computers and Mathematics with Applications, 66 (2013), pp. 1058–1064.
- [27] ———, *Discontinuous Petrov Galerkin (DPG) method with optimal test functions*. ECCOMAS Newsletter, European Community on Computational Methods in Applied Sciences, December 2014.
- [28] L. DEMKOWICZ AND J. LI, *Numerical simulations of cloaking problems using a DPG method*, Comput. Mech., 51 (2013), pp. 661–672.
- [29] T. ELLIS, L. DEMKOWICZ, AND J. CHAN, *Locally conservative discontinuous Petrov-Galerkin finite elements for fluid problems*, Comput. Math. Appl., 68 (2014), pp. 1530–1549.

- [30] T. E. ELLIS, L. F. DEMKOWICZ, J. L. CHAN, AND R. D. MOSER, *Space-time DPG: Designing a method for massively parallel CFD*, Comput. & Fluids, (2015 (to appear)).
- [31] I. EPANOMERITAKIS, V. AKÇELİK, O. GHATTAS, AND J. BIELAK, *A Newton-CG method for large-scale three-dimensional elastic full-waveform seismic inversion*, Inverse Problems, 24 (2008), p. 034015 (26pp).
- [32] F. FUENTES, B. KEITH, L. DEMKOWICZ, AND S. NAGARAJ, *Orientation embedded high order shape functions for the exact sequence elements of all shapes*, Comput. Math. Appl., 70 (2015), pp. 353–458.
- [33] A. GHOLAMI, D. MALHOTRA, AND G. BIROS, *AccFFT: A library for distributed-memory FFT on CPU and GPU architectures*, SIAM Journal on Scientific Computing, (2016), pp. 1–15. In print.
- [34] A. GHOLAMI, D. MALHOTRA, H. SUNDARY, AND G. BIROS, *FFT, FMM, or multigrid? a comparative study of state-of-the-art Poisson solvers*, SIAM Journal on Scientific Computing, (2016).
- [35] J. GOPALAKRISHNAN, S. LANTERI, N. OLIVARES, AND R. PERRUSSEL, *Stabilization in relation to wavenumber in HDG methods*, Advanced Modeling and Simulation in Engineering Sciences, 2 (2015), p. Article 13.
- [36] J. GOPALAKRISHNAN, F. LI, N.-C. NGUYEN, AND J. PERAIRE, *Spectral approximations by the HDG method*, Math. Comp., 84 (2014), pp. 1037–1059.
- [37] J. GOPALAKRISHNAN, P. MONK, AND P. SEPÚLVEDA, *A tent pitching scheme motivated by Friedrichs theory*, Computers and Mathematics with Applications, 70 (2015), pp. 1114–1135.
- [38] J. GOPALAKRISHNAN, I. MUGA, AND N. OLIVARES, *Dispersive and dissipative errors in the DPG method with scaled norms for the Helmholtz equation*, SIAM J. Sci. Comput., 36 (2014), pp. A20–A39.
- [39] J. GOPALAKRISHNAN AND W. QIU, *An analysis of the practical DPG method*, Mathematics of Computation, 83 (2014), pp. 537–552.
- [40] J. GOPALAKRISHNAN AND J. SCHÖBERL, *Degree and wavenumber [in]dependence of Schwarz preconditioner for the DPG method*, in Spectral and High Order Methods for Partial Differential Equations, R. M. Kirby, M. Berzins, and J. S. Hesthaven, eds., Springer, 2015, pp. 257–265. Selected papers from the ICOSAHOM conference, June 23–27, 2014, Salt Lake City, Utah, USA.
- [41] —, *Software hosted on GitHub: DPG Methods in NGSolve*, 2015. <https://github.com/jayggg/DPG>.
- [42] B. KEITH, F. FUENTES, AND L. DEMKOWICZ, *The DPG methodology applied to different variational formulations of linear elasticity*, tech. rep., ICES, January 16–01. submitted to CMAME.
- [43] D. MALHOTRA AND G. BIROS, *PVFMM: A distributed memory fast multipole method for volume potentials*, ACM Transactions on Mathematical Software, (2016). To appear.
- [44] D. MALHOTRA, A. GHOLAMI, AND G. BIROS, *A volume integral equation Stokes solver for problems with variable coefficients*, in Proceedings of SC14, The SCxy Conference series, New Orleans, Louisiana, November 2014, ACM/IEEE. Acceptance rate 82/394 (21%), nominated for best student paper award.

- [45] W. B. MARCH AND G. BIROS, *Far-field compression for fast kernel summation methods in high dimensions*, Applied and Computational Harmonic Analysis, (2015).
- [46] W. B. MARCH, B. XIAO, AND G. BIROS, *ASKIT: Approximate skeletonization kernel-independent treecode in high dimensions*, SIAM Journal on Scientific Computing, 37 (2015), pp. 1089–1110.
- [47] W. B. MARCH, B. XIAO, S. THARAKAN, C. D. YU, AND G. BIROS, *A kernel-independent FMM in general dimensions*, in Proceedings of SC15, The SCxy Conference series, Austin, Texas, November 2015, ACM/IEEE. Acceptance rate 79/358 (22%).
- [48] ———, *Robust treecode approximation for kernel machines*, in Proceedings of the 21st ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Sydney, Australia, August 2015, pp. 1–10. 20% acceptance, 159/819.
- [49] W. B. MARCH, B. XIAO, C. YU, AND G. BIROS, *An algebraic parallel treecode in arbitrary dimensions*, in Proceedings of IPDPS 2015, 29th IEEE International Parallel and Distributed Computing Symposium, Hyderabad, India, May 2015. acceptance rate 107/495 (21.8%).
- [50] W. B. MARCH, B. XIAO, C. D. YU, AND G. BIROS, *ASKIT: An efficient, parallel library for high-dimensional kernel summations*, SIAM Journal on Scientific Computing, (2016). to appear.
- [51] J. MARTIN, L. C. WILCOX, C. BURSTEDDE, AND O. GHATTAS, *A stochastic Newton MCMC method for large-scale statistical inverse problems with application to seismic inversion*, SIAM Journal on Scientific Computing, 34 (2012), pp. A1460–A1487.
- [52] P. J. MATUSZYK AND L. F. DEMKOWICZ, *Parametric finite elements, exact sequences and perfectly matched layers*, Comput. Mech., 51 (2013), pp. 35–45.
- [53] N. PETRA, J. MARTIN, G. STADLER, AND O. GHATTAS, *A computational framework for infinite-dimensional Bayesian inverse problems: Part II. Stochastic Newton MCMC with application to ice sheet inverse problems*, SIAM Journal on Scientific Computing, 36 (2014), pp. A1525–A1555.
- [54] J. POULSON, B. MARKER, R. A. VAN DE GEIJN, J. R. HAMMOND, AND N. A. ROMERO, *Elemental: A new framework for distributed memory dense matrix computations*, ACM Transactions on Mathematical Software (TOMS), 39 (2013), p. 13.
- [55] B. QUAIFE AND G. BIROS, *On preconditioners for the Laplace double-layer in 2d*, Numerical Linear Algebra with Applications, (2014).
- [56] N. V. ROBERTS, *Camellia: A software framework for discontinuous Petrov-Galerkin methods*, Computers & Mathematics with Applications, 68 (2014), pp. 1581 – 1604.
- [57] J. RUDI, A. C. I. MALOSSO, T. ISAAC, G. STADLER, M. GURNIS, Y. INEICHEN, C. BEKAS, A. CURIONI, AND O. GHATTAS, *An extreme-scale implicit solver for complex PDEs: Highly heterogeneous flow in earth's mantle*, in SC15: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, ACM, 2015, pp. 5:1–5:12. Winner of Gordon Bell Prize.
- [58] S. P. S. NAGARAJ AND L. DEMKOWICZ, *Construction of DPG Fortin operators for second order problems*, Tech. Rep. 22, ICES, 2015.

- [59] H. SUNDAR, G. BIROS, C. BURSTEDDE, J. RUDI, O. GHATTAS, AND G. STADLER, *Parallel geometric-algebraic multigrid on unstructured forests of octrees*, in SC12: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, Salt Lake City, UT, 2012, ACM/IEEE.
- [60] H. SUNDAR AND O. GHATTAS, *A nested partitioning algorithm for adaptive meshes on heterogeneous clusters*, in ACM International Conference on Supercomputing, ICS'15, Newport Beach, CA, 2015.
- [61] H. SUNDAR, G. STADLER, AND G. BIROS, *Comparison of multigrid algorithms for high-order continuous finite element discretizations*, Numerical Linear Algebra with Applications, 22 (2015), pp. 664–680.
- [62] ———, *Comparison of multigrid algorithms for high-order continuous finite element discretizations*, Numerical Linear Algebra with Applications, 22 (2015), pp. 664–680.
- [63] K. WANG, T. BUI-THANH, AND O. GHATTAS, *A randomized maximum a posteriori method for posterior sampling of high dimensional nonlinear Bayesian inverse problems*. Submitted, 2016.
- [64] L. C. WILCOX, G. STADLER, T. BUI-THANH, AND O. GHATTAS, *Discretely exact derivatives for hyperbolic PDE-constrained optimization problems discretized by the discontinuous Galerkin method*, Journal of Scientific Computing, 63 (2015), pp. 138–162.
- [65] C. YU, W. MARCH, B. XIAO, AND G. BIROS, *INV-ASKIT: a parallel fast direct solver for kernel matrices*, in 30th IEEE International Parallel & Distributed Processing Symposium (IEEE IPDPS 2016), Chicago, USA, May 2016.
- [66] H. ZHU, S. LI, S. FOMEL, G. STADLER, AND O. GHATTAS, *A Bayesian approach to estimate uncertainty for full waveform inversion with a priori information from depth migration*. Submitted, 2015.

1.

**1. Report Type**

Final Report

**Primary Contact E-mail****Contact email if there is a problem with the report.**

omar@ices.utexas.edu

**Primary Contact Phone Number****Contact phone number if there is a problem with the report**

5129499818

**Organization / Institution name**

The University of Texas at Austin

**Grant/Contract Title****The full title of the funded effort.**

Ultra-Scalable Algorithms for Large-Scale Uncertainty Quantification in Inverse Wave Propagation

**Grant/Contract Number****AFOSR assigned control number. It must begin with "FA9550" or "F49620" or "FA2386".**

FA9550-12-1-0484

**Principal Investigator Name****The full name of the principal investigator on the grant or contract.**

Omar Ghattas

**Program Manager****The AFOSR Program Manager currently assigned to the award**

Jean-Luc Cambier

**Reporting Period Start Date**

09/30/2012

**Reporting Period End Date**

11/30/2015

**Abstract**

The overall aim of this project was to develop scalable algorithms for the inverse problem of inferring, with associated uncertainty, the heterogeneity of a medium or shape of a scatterer from reflected/transmitted waves (acoustic, elastic, electromagnetic) at very large scale. The resulting Bayesian wave inverse propagation problem has been intractable using contemporary algorithms. Research was conducted under three complementary subprojects. The first subproject (led by O. Ghattas) focused on scalable algorithms for large-scale Bayesian inverse problems governed by time domain wave propagation. The second subproject (led by G. Biros) focused on fast algorithms for inverse scattering and uncertainty quantification based on volume integral equation formulations for the inverse medium problem. The third subproject (led by L. Demkowicz and J. Gopalakrishnan) focused on new, highly efficient discretizations for wave propagation in the form of the discontinuous Petrov Galerkin (DPG) method and associated solvers. Results and conclusions in each sub-project area are discussed in separate sections of the report.

**Distribution Statement****This is block 12 on the SF298 form.**

Distribution A - Approved for Public Release

**Explanation for Distribution Statement**

DISTRIBUTION A: Distribution approved for public release.



If this is not approved for public release, please provide a short explanation. E.g., contains proprietary information.

#### **SF298 Form**

Please attach your [SF298](#) form. A blank SF298 can be found [here](#). Please do not password protect or secure the PDF  
The maximum file size for an SF298 is 50MB.

[SF298.pdf](#)

**Upload the Report Document.** File must be a PDF. Please do not password protect or secure the PDF . The maximum file size for the Report Document is 50MB.

[afosr-final-report-feb16.pdf](#)

**Upload a Report Document, if any.** The maximum file size for the Report Document is 50MB.

**Archival Publications (published) during reporting period:**

See list of publications in the report.

**Changes in research objectives (if any):**

**Change in AFOSR Program Manager, if any:**

AFOSR Computational Mathematics Program Manager Dr. Jean Luc Cambier replaced Dr. Fariba Fahroo during the last 6 months of the grant.

**Extensions granted or milestones slipped, if any:**

**AFOSR LRIR Number**

**LRIR Title**

**Reporting Period**

**Laboratory Task Manager**

**Program Officer**

**Research Objectives**

**Technical Summary**

**Funding Summary by Cost Category (by FY, \$K)**

	Starting FY	FY+1	FY+2
Salary			
Equipment/Facilities			
Supplies			
Total			

**Report Document**

**Report Document - Text Analysis**

**Report Document - Text Analysis**

**Appendix Documents**

**2. Thank You**

**E-mail user**

Feb 29, 2016 18:46:41 Success: Email Sent to: omar@ices.utexas.edu